The Shield that Never Was: Societies with Single-Peaked Preferences are More Open to Manipulation and Control[*]

Piotr Faliszewski
Department of Computer Science
AGH University of Science and Technology
30-059 Kraków, Poland

Edith Hemaspaandra
Department of Computer Science
Rochester Institute of Technology
Rochester, NY 14623, USA

Lane A. Hemaspaandra
Department of Computer Science
University of Rochester
Rochester, NY 14627, USA

Jörg Rothe
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
40225 Düsseldorf, Germany

September 17, 2009; revised June 19, 2010

Abstract

Much work has been devoted, during the past twenty years, to using complexity to protect elections from manipulation and control. Many results have been obtained showing NP-hardness shields, and recently there has been much focus on whether such worst-case hardness protections can be bypassed by frequently correct heuristics or by approximations. This paper takes a very different approach: We argue that when electorates follow the canonical political science model of societal preferences the complexity shield never existed in the first place. In particular, we show that for electorates having single-peaked preferences, many existing NP-hardness results on manipulation and control evaporate.

1 Introduction

Elections are a broadly used preference aggregation model in both human societies and multiagent systems. For example, elections have been proposed as a mechanism for collaborative

1

decision-making in such multiagent system contexts as recommender systems/collaborative filtering [PHG00] and planning [ER91,ER97]. The importance of elections explains why elections are studied intensely over a wide range of fields, including political science, mathematics, social choice, artificial intelligence, economics, and operations research. Formally, an election system takes as input a set of candidates (or alternatives) and a set of votes (usually each a 0-1 approval vector over the candidates or a linear order over the candidates) and outputs a subset of the candidates as the winner(s).

Control refers to attempts to make a given candidate win [BTT92] (or not win [HHR07]) an election by such participation-change actions as adding or deleting voters or candidates. Manipulation refers to attempts to make a given candidate win [BTT89,BO91] (or not win [CSL07]) by some coalition of voters who strategically change their votes. There is a large literature, started by the insightful contributions of Bartholdi, Orlin, Tovey, and Trick [BO91,BTT89,BTT92], on choosing election systems that make control and manipulation NP-hard, i.e., on choosing election systems that seek to make control and manipulation computationally prohibitive (see the survey [FHHR09b]). Recently, there has been a flurry of work seeking to bypass worst-case manipulation hardness results by frequently correct heuristics or approximation algorithms (as a few pointers into that literature, see, e.g., [FKN08,XC08a,XC08b,CS06,PR07,BFH+08]).

The present paper takes a radically different approach. We study elections where the vote set must be "single-peaked." We'll discuss single-peakedness in more detail after stating our main contributions. But, briefly put, single-peakedness means that there is some linear ordering of the candidates relative to which (in the model in which votes are linear orders) each voter's preferences always increase, always decrease, or first increase and then decrease, or (in the model in which votes are approval vectors) each voter's approved candidates are contiguous within the linear order. Single-peaked preferences, introduced by Black [Bla48] (see also [Bla58]), model societies that are heavily focused on one issue (taxes, war, etc.), and the single-peaked framework is so central to political science that it has been described as "*the* canonical setting for models of political institutions" [GPP], and indeed is typically the model of societal voting first covered in an introductory course on positive (i.e., theoretical) political science. This paper's main contributions are the following:

1. We introduce the study of single-peakedness for approval-voting elections.

2. In Section 3 we show that for both voting by approval vectors and for voting by linear orders, many election control problems known to be NP-hard in the general case have polynomial-time algorithms in the single-peaked case.

3. In Section 4 we show that many election manipulation problems known to be NP-hard in the general case have polynomial-time algorithms in the single-peaked case.

   However—and in this we are inspired by the path-setting work of Walsh [Wal07], who showed that Single Transferable Vote (for at least 3 candidates and weighted votes) remains NP-hard to manipulate even in the single-peaked case—we in Section 4 also show that many manipulation problems remain NP-hard even when restricted to the single-peaked case.

2

4. We show, contrary to intuition and the tacit assumptions of some papers, that even for natural systems there are cases where increasing the number of candidates decreases the complexity. In particular, we as Theorem 4.2 show that in the single-peaked case in 3-veto elections (i.e., one votes against three candidates and for all others) manipulation is in P for up to four candidates, is NP-hard for five candidates, and is in P for six or more candidates.

We mostly defer discussion of related work until after our results, as the related work will then have more context and definitions to draw on. However, we mention now a point raised by a conference referee: Since "median voting" and "generalized median voter schemes" are known to in many settings never give a voter an incentive to vote insincerely for single-peaked electorates (see [Bar01] and [PS99], and the references therein, for definitions and discussion), shouldn't single-peaked electorates use only median voting? Our reply is that in real-life scenarios voting rules are typically fixed (e.g., to be plurality or approval) and cannot be changed when one suspects the electorate to be single-peaked. In addition, lack of incentive to vote insincerely regarding manipulation does not say anything about a rule's resistance to control. Thus, we believe that our results are very interesting and relevant.

2 Preliminaries

Elections and Preferences. An *election* consists of a set C of candidates and a collection V of votes. We will consider two different models for votes. One is that each vote is a vector (an *approval vector*) from $\{0,1\}^{\|C\|}$, denoting approval (1) or disapproval (0) for each candidate. The other is that each vote is a linear order (by which we mean a strict, linear order—a complete, transitive, antisymmetric relation) over the candidates, e.g., Bob > Alice > David > Carol. An *election system* is a mapping that takes as input a candidate set C and a set V of votes over that candidate set, and outputs an element of 2^C, i.e., outputs which candidates are *winners* of the election. (We, like Bartholdi, Tovey, and Trick [BTT92], do not expressly forbid elections with no winners, although all of the many natural election systems discussed in this paper have the property that there is always at least one winner when there is at least one candidate.) Except where we explicitly state otherwise, V is a list of votes (ballots), so if three votes are the same, they will appear three times in the list. We will use *succinct input* [FHH09] to describe the quite different input model in which each preference that is held by one or more voters appears just once in the list and is accompanied by a binary number stating how many voters have that preference.

The election systems of most interest to us in this paper are the following ones. In *approval voting*, voters vote by approval vectors, and whichever candidate(s) get the most approvals are the winner(s). A *scoring protocol* election, which is always defined for a specific number m of candidates, is specified by a *scoring vector*, $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_m) \in \mathbb{N}^m$, satisfying $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_m$. Votes are linear orders. Each vote contributes α_1 points to that vote's most preferred candidate, α_2 points to that vote's second most preferred candidate, and so on. And whichever candidate(s) get the most total points are the winner(s). Among the most important scoring protocols for m candidates are plurality,

with $\alpha = (1, \overbrace{0, \ldots, 0}^{m-1})$, j-veto ($j \leq m$), with $\alpha = (\overbrace{1, \ldots, 1}^{m-j}, \overbrace{0, \ldots, 0}^{j})$, and Borda, with $\alpha = (m - 1, m - 2, \ldots, 0)$. For m candidates, j-approval and $(m - j)$-veto are the same system. We'll also speak of the veto system for an unbounded number of candidates, by which we mean that on inputs with m candidates, each voter gives 1 point to all candidates other than her least favorite candidate and 0 points to her least favorite candidate. We'll similarly also speak of plurality for an unbounded number of candidates, by which we mean that on inputs with m candidates, each voter gives 0 points to all candidates other than her favorite candidate and 1 point to her favorite candidate. And the general case of approval voting is always for an unbounded number of candidates.

Let us fix some notation for the election systems defined above. Let (C, V) be an election and let c be a candidate in C. By $score_{(C,V)}(c)$ we mean the score of c in election (C, V), i.e., the number of points c receives under plurality, veto, or some given scoring protocol and the number of c's approvals under approval voting (it will always be clear from the context, which election system the score refers to). If the candidate set C is clear from the context, we simply write $score_V(c)$.

Single-Peaked Preferences. A collection V of votes, each vote v_i being a linear order $>_i$ over C, is said to be *single-peaked* exactly if there exists a linear order over C, call it L, such that for each triple of candidates c, d, and e, it holds that:

$$(c \, L \, d \, L \, e \vee e \, L \, d \, L \, c) \Longrightarrow (\forall i) \, [c >_i d \Longrightarrow d >_i e].$$

That is just a formal way of saying that with respect to L, each voter's degree of preference rises to a peak and then falls (or just rises or just falls). The loose intuition behind this is captured in Figure 1. If we imagine an electorate completely focused on one issue (say taxes), with each person having a single-peaked (in the natural analogous sense of that notion applied to curves) utility curve (but potentially different utility curves for different voters), one gets precisely this notion (give or take the issue of ties). In Figure 1's example, the preferences of v_1 would be $c_1 > c_2 > c_3 > c_4 > c_5$, of v_2 would be $c_3 > c_4 > c_2 > c_1 > c_5$, and of v_3 would be $c_4 > c_3 > c_2 > c_1 > c_5$.

The seminal study of single-peaked preferences was done by Black [Bla48], and that work and many subsequent studies (e.g., [NW87, DHO70, PR97, Kre98, Bla58]) argued that single-peaked preferences (in the unidimensional spatial model) are a broadly useful model of electoral preferences that captures many important settings.[1] Of course, issues that are multidimensional—as many issues are—are typically not captured by single-peaked preferences (by which we always mean the unidimensional case). And even in a society that is completely focused on one issue, some maverick individuals may focus on other issues, so one should keep in mind that the single-peaked case is a widely studied but extreme model.

---

[1]We mention in passing that since with single-peaked preferences there is always a Condorcet winner when the number of voters is odd, every Condorcet-consistent voting rule has an efficient winner-determination algorithm for odd numbers of voters [Wal07]. Of course, many interesting voting rules are not Condorcet-consistent.

Figure 1: Single-Peaked Preference Motivation: Utility Curves for Three Voters.

Looking again at Figure 1, let us imagine that each voter has a utility threshold at which she starts approving of candidates. Then in the approval-vector-as-method-of-voting we'll have that for some linear order L, each voter either disapproves of everyone or approves of precisely a set of candidates that are contiguous with respect to L. Formally, we say a collection V, made up of approval vectors $v_1, v_2, \ldots, v_n$ over the candidate set C, with Approves$_i$ being the set of candidates that v_i approves, is *single-peaked* exactly if there is a linear order over C, call it L, such that for each triple of candidates c, d, and e, it holds that:

$$c\,L\,d\,L\,e \implies (\forall i)\,[\{c, e\} \subseteq \text{Approves}_i \implies d \in \text{Approves}_i].$$

The reasons that single-peaked approval voting is compellingly natural to study (and the reasons why one should not go overboard and claim that it is a universally appropriate notion) are essentially the same as those, touched on above, involving single-peaked voting with respect to preference orders. In many cases it is natural to assume that there is some unidimensional issue steering society and that each person's range of comfort on that issue is some contiguous segment along that issue's dimension. In such a case, each person's set of approved candidates will form a contiguous block among the candidates when they are ordered by their positions on that issue. To the best of our knowledge, this is the first paper to study elections based on single-peaked approval vectors—although the literature is so vast that we would not be at all surprised such elections had been previously studied.

In our control and manipulation problems, we'll for the single-peaked case follow Walsh [Wal07] and take as part of the input a particular linear order of the candidates relative to which the votes are single-peaked. This is arguably natural—as we may view candidates' positions on the issue that defines the entire election as being openly known. However, one may wonder how hard it is to, given a set of voters, tell whether it is single-peaked. For the case of votes being linear orders, Bartholdi and Trick ([BT86], see also [DF94, ELÖ08]; a

5

somewhat related paper is [Tri89]) show by a path-based graph algorithm that the problem is in P, and Doignon and Falmagne [DF94] and Escoffier, Lang, and Öztürk [ELÖ08] show how to produce in polynomial time a linear order witnessing the single-peakedness when such an ordering exists. For the case of votes being approval vectors, the literature already has long contained the analogue of both the results just mentioned. In particular, the work of Fulkerson and Gross [FG65, Sections 5 and 6] and Booth and Lueker [BL76, Theorem 6] proves (rephrased in our terminology) that in polynomial time—in a certain natural sense, even in linear time—one can determine whether a set of approval vectors is single-peaked, can in such a case efficiently find a linear ordering realizing the single-peakedness, and indeed can implicitly represent all linear orderings realizing the single-peakedness.

Theorem 2.1 ([FG65, BL76]) *Given a collection V of approval vectors over C, in polynomial time we can produce a linear order L witnessing V's single-peakedness or can determine that V is not single-peaked.*

Note that Theorem 2.1 does not seem to follow from the analogous result for linear orders. For example, extending each approval vector to form a linear order with all approved candidates preceding all the disapproved ones, and then running the algorithm of Escoffier, Lang, and Öztürk [ELÖ08] on the thus transformed input does not work.

Control and Manipulation Problems. For an election system $\mathcal{E}$, the *Constructive Control by Adding Candidates* problem is the set of all (C, V, p, k, C'), where V consists of votes over $C \cup C'$, $p \in C$, and $C \cap C' = \emptyset$, such that there is a set $C'' \subseteq C'$ with $\|C''\| \leq k$ for which candidate p is the unique winner under election system $\mathcal{E}$ when the voters in V vote over $C \cup C''$ (i.e., we restrict each voter down to her induced preferences over $C \cup C''$). The *Destructive Control by Adding Candidates* problem is the same except now the question is whether there is such a C'' ensuring that p is *not* a unique winner.

The Constructive/Destructive Control by Deleting Candidates problems are analogous, with inputs of the form (C, V, p, k), where k is a limit on how many candidates from C can be deleted, and it is forbidden to delete p. The Constructive/Destructive Control by Adding/Deleting Voters problems are analogous, with inputs respectively (C, V, p, k, V') and (C, V, p, k), where V', $V' \cap V = \emptyset$, is a pool of "unregistered" voters, and k denotes the bound on how many voters from V' we can add (Adding case) or how many voters from V we can delete (Deleting case).

Constructive/Destructive Control by Unlimited Adding Candidates is the same as Constructive/Destructive Control by Adding Candidates except there is no "k"; one may add any or all of the members of C'.

Most of these problems were introduced in the seminal control paper of Bartholdi, Tovey, and Trick [BTT92], and the remaining ones were introduced in [HHR07, FHHR09a]. These problems model such real-world problems as introducing a candidate to run to split off another candidate's support; urging an independent candidate to withdraw for the good of the country; spreading rumors that people with outstanding warrants who try to vote will

be arrested; and sending vans to retirement homes to drive car-less members of one's party to the voting place.

In each case, the inputs we specified are for the "general" case. For the single-peaked case there will be an additional input, a linear order L over the candidates such that relative to L the election is single-peaked (with respect to all voters and candidates—even those in C' and in V' in problems having those extra sets). (If L is not such a linear order, the input is immediately rejected. L's goodness can be easily tested in polynomial time, simply by looking at each vote.) As mentioned previously, in assuming that L is given we are following Walsh's [Wal07] natural model—L is the broadly known positioning of the candidates.[2]

The (constructive coalition weighted) manipulation problem was introduced by Conitzer, Sandholm, and Lang ([CSL07] and its conference precursors, building on the seminal manipulation papers [BO91,BTT89]) and takes as input a list of candidates C, a list of non-manipulative voters each specified by preferences (as a linear order or as an approval vector, depending on our election system) over C and a nonnegative integer weight, a list of the weights of the voters in our manipulating coalition, and the candidate p that our coalition seeks to make a winner. The set of all such inputs for which there is some assignment of preferences to the manipulators that makes p a winner is the *Constructive Coalition Weighted Manipulation* problem for that election system. The *Constructive Size-k-Coalition Unweighted Manipulation* problem is defined analogously, except with all voters having unit weight and with k being the number of manipulators.

We follow the convention from [BTT89,BO91,BTT92] of focusing on the unique-winner case for control and the winner case (i.e., asking whether the candidate can be made *a* winner or be prevented from being *a* winner; following the literature, we'll sometimes refer to that as the "nonunique-winner model/case") for manipulation. In many cases we've shown a given result in both models and will sometimes mention that in passing. All of the results of Conitzer, Sandholm, and Lang of relevance to us hold in both models (see Footnote 7 of [CSL07]).

For the single-peaked case, a linear order L is given as part of the input and the manipulating coalition's votes must all be single-peaked with respect to L (as must all other voters) for the input to be accepted.

Complexity Notions for Control and Manipulation Problems. If by a given type of constructive control action we can never change p from not unique winning (not winning)

---

[2]Both for linear-order votes ([DF94], see also [BT86,ELÖ08]) and for approval-vector votes (Theorem 2.1) it holds that even if L is not given one can find a good L in polynomial time if one exists. This fact may be comforting to those who would prefer that L not be given. But we caution that a given vote set V may have many valid L's. And for some problems, which valid L one uses may affect the problem and its complexity. Indeed, not having L as part of the input might even open the door to time-sequence issues, e.g., in deletion of candidates/voters, should we ask instead whether single-peakedness should only have to hold after the deletion? (We'd say, ideally, "no.") These issues are reasonable for further control/manipulation-related study (and the issue of whether L is known *a priori* is much studied in the political economy literature already, see, e.g., [AB04, Section 2.4]), but we find the Walsh [Wal07] model to be most natural and compelling. In many cases our proofs work fine if L is not part of the input, and at times we'll mention that.

to unique winning (winning) we say the problem is *immune* to that control type, and otherwise we call it *susceptible* to that control type. The destructive case is analogous. If a susceptible problem is in P we call it *vulnerable*, and if it is NP-hard we call it *resistant*. A manipulation problem (such as the constructive coalition weighted manipulation problem defined above) is said to be *vulnerable* if it is in P and *resistant* if it is NP-hard. In each control/manipulation case in this paper where we assert vulnerability (or membership in P), not only is the decision problem in P but also in polynomial time we can produce a successful control/manipulation action if one exists (i.e., the problem is what [HHR07] calls *certifiably vulnerable*). Most of these notions are in detail from or in the general spirit of the work of Bartholdi, Orlin, Tovey, and Trick [BTT92,BTT89,BO91], as in some cases naturally modified or extended in [HHR09,HHR07].

Almost all NP-hardness proofs of this paper follow via reductions from the well-known NP-complete problem PARTITION (see, e.g., [GJ79]). Specifically, we use the variant of the problem where our input contains a set $\{k_1, \ldots, k_n\}$ of n distinct positive integers that sum to $2K$, and we ask: Does there exist a subset $A \subseteq \{1, \ldots, n\}$ such that $\sum_{i \in A} k_i = K$?

3 Control

3.1 Results

We now turn to our results on control of elections. The theme of this paper is that electorates limited to being single-peaked often are simpler to control and manipulate. Intuitively speaking, we will show that the more limited range of vote collections allowed by single-peaked voting (as opposed to general voting) is so restrictive that the reductions showing NP-hardness fall apart—that those reductions centrally need complex collections of votes to work. Of course, for all we know, perhaps P = NP, and so attempting to show that no reduction at all can exist from SAT to our single-peaked control problems would be a fool's errand (or a Turing Award-scale quest). Rather, we'll show our single-peaked control problems easy the old-fashioned way: We'll prove that they are in P. The techniques we use to prove them in P vary from easy observations to smart greedy schemes to dynamic programming.

The control complexity of approval voting is studied in detail by Hemaspaandra, Hemaspaandra, and Rothe [HHR07] (see [ENR09] regarding limited adding of candidates, see also the survey [BEH$^+$] for an overview of computational properties of approval voting). Among all the adding/deleting control cases defined in this paper (ten in total), precisely two are resistant, and the other eight (all five destructive cases and the three constructive candidate cases) are immune or vulnerable. The two resistant cases are Constructive Control by Adding Voters and Constructive Control by Deleting Voters.

The following theorem shows that both of these complexity shields evaporate for societies with single-peaked preferences. Briefly put, the challenge here is that the set V' of voters to add may be filled with "incomparable" voter pairs—pairs such that regarding their interval with respect to the linear order defining single-peakedness, neither is a subset of the other—

and so it is not immediately obvious what voters to add. We solve this by a "smart greedy" approach, breaking votes first into broad groups based on where their intervals' right edges fall with respect to just a certain "dangerous" subset of the candidates, and then re-sort those based on their left edges, and we argue that if any strategy will reach the control goal then this one will.

Theorem 3.1 *For the single-peaked case, approval voting is vulnerable to constructive control by adding voters and to constructive control by deleting voters, in both the unique-winner model and the nonunique-winner model, for both the standard input model and the succinct input model.*[3]

Proof.   We'll cover the unique-winner cases only. From these it will be completely clear how to handle the nonunique-winner cases. We'll cover just the succinct-input cases, as each implies the corresponding standard-input case. Susceptibility is immediately clear by easy examples.

First, let's speak of the societal linear order. In the model in which the (purported) linear order is part of the input, start by doing the obvious polynomial-time check to see whether this linear order truly is valid with respect to the voters V, the unregistered voters V', and the candidates C. That is, check to ensure that each member of V and V' approves either of no candidates or approves of a exactly a collection of candidates that are contiguous with respect to the input linear order. If so, the votes of $V \cup V'$, which are approval vectors over C, are single-peaked with respect to the input societal linear order over C. If not, output "Invalid societal linear order" and halt. (This very easy observation in effect shows that, for the approval-vector case, checking a given purported societal linear order is in P. That is equally obviously true for the votes-are-linear-orders setting, of course.)

If we're in the model in which a linear order must be found, if one exists, relative to which the problem is single-peaked and the given control action is possible, simply use Theorem 2.1 to efficiently find some linear order relative to which our input is single-peaked. If none exists, state that and terminate. If Theorem 2.1 gives us some valid order, it is not hard to see that the algorithm we are going to give in this proof will find a successful control action if and only if there exists any valid societal linear order relative to which a successful control action exists. That is basically because if π_1 and π_2 are valid societal linear orders, the votes are the same either way, and in each case the algorithm we are about to give will succeed exactly if there exists a legal-cardinality collection of unregistered voters that achieves the control goal.

So, let us take it that we have the societal linear order, which we'll refer to as L. We start with the case of adding voters. This proof will be done rather visually, and although we'll implicitly be sketching a general proof, we'll keep close to heart an example, as that will make the idea of the proof clearer. Let p be the preferred candidate—the one we seek to

---

[3]This result holds in our settled model in which the linear order specifying the society's order on the candidates is part of the input, and also holds in the model in which the linear order is not part of the input but rather the question is whether there exists any valid linear order relative to which there is a way of achieving our control goal.

Figure 2: Bar chart of approvals from V, plus intervals indicating the votes in V' and their multiplicities. The candidates are ordered in the societal linear order and the "dangerous" candidates are marked with shaded bars.

make a unique winner. Figure 2 shows a representation of the initial state. In this figure, we have (for display purposes) renamed all the candidates in a way that echos the societal linear order L, i.e., the candidates are each labeled with respect to whether they are left or right of p in L, and are ordered following the societal order. The bars represent the total number of approvals each candidate gets over all the votes contained in V. (Since we're in the succinct-input case, we of course do in computing that take into account the multiplicities, e.g., if candidate r_{10} is approved of in one ordering that occurs with multiplicity 4 and one that occurs with multiplicity 3, and in no others, then the bar for r_{10} would show that it gets 7 (not 2) approvals.) The intervals below the bar graph reflect the votes in V'. Note that each vote type in V' either approves of no candidates (in which case we won't even display it; it is of no help toward meeting our control goal) or, necessarily, is an interval (although perhaps a degenerate one that contains exactly one candidate) with respect to L. (Since we're handling the succinct case, votes of the same vote type are in the figure listed together along with their multiplicity to their right. We do this because in the input to the problem votes come in by type and with a listed-in-binary multiplicity, in the succinct case.)

10

The number of vote types is $O(\|C\|^2)$, since each vote type (except for disapproving of everyone) is defined by its left and right endpoints. And, even more to the point, the number of vote types is trivially at most linear in the size of the input, since each vote type cast by at least one voter has to appear in the input. So the number of vote types certainly is not so large as to cause any challenges (about having enough time to consider the vote types occurring in V') regarding remaining within polynomial time.

Now, at first, things might seem worrisome. How can we decide which votes to add, especially between "incomparable" vote types such as (in our figure) the vote type with multiplicity 2 and the vote type with multiplicity 7? Those two vote types each help p and some other candidates, but the sets of other candidates helped are incomparable—neither is a subset of the other.

We handle this worry by giving a "smart greedy" algorithm that only makes moves that are clearly at least as safe as any other move, e.g., that will never close down the last path to success if any path to success exists.

Let us look again at the figure. Note that although there are many candidates, we can ignore all candidates except p herself and the candidates that are shaded—l_6, l_3, l_2, r_3, r_6, and r_{10}. Why? Well, in our algorithm, we will never add any vote from V' that does not approve of p. In fact, let us from now on consider V' to be redefined to have removed from it every vote that does not approve of p. (So in the figure, the V' vote types shown with multiplicity 5 and 9 go, as do any votes inside "etc." that do not approve of p.) Now, since we'll never add a vote from V' that does not approve of p (indeed, we redefined V' to drop such votes), and since each added vote will be an interval that includes p, notice that, for example, if we add some votes from V' that cause p to have (in total, between the original votes from V and added votes from V') strictly more approvals than r_3, then p necessarily will have strictly more approvals than r_1 and r_2, since each will have received from the added votes from V' no more additional approvals than p received from those added votes. By the same reasoning, if by adding votes from (our redefined) V' we ensure that p has strictly more approvals than r_6, clearly p will have strictly more approvals than r_4 and r_5.[4] And the same argument applies to the case of r_{10} and, in mirror image, the shaded left-side bars. So, these are the only "dangerous" contenders that we have to worry about regarding p. This observation will guide our choice of which votes to add and which candidates to ignore.

Before going on, let us mention why we in the above went step-by-step. That is, one might ask why we did not simply say that if p by adding votes from (our redefined) V' could

---

[4]The reader may worry that this "by the same reasoning" is not a valid claim, since one may worry that since r_5 starts above p in score, we may by adding votes that approve of both p and r_5 (but not r_6) end up with p having strictly more approvals than r_6 and yet with r_5 having more approvals than p, contrary to what our sentence above is asserting. The reason this reasonable worry is not actually a problem is that we will (as will be made clearer later in the proof) use a particular step-by-step process. The process we follow will, before even thinking of adding any votes approving r_5, first have (if there is any path to ensuring this; if not, control is impossible) ensured *via adding votes that give approval to p but not to r_3 (and thus not to r_5)* that p has more approvals than r_3. But note that after doing that, our "by the same reasoning" does cleanly apply to the case of r_4, r_5, and r_6. This point is not identical to, but rather works hand-in-hand with, the point of the paragraph following the current paragraph.

be made to defeat r_{10} then it necessarily would defeat all of r_1 through r_9. But that claim in fact does not hold. Consider the case of adding 6 of the vote type having multiplicity 7 in the figure. Doing so boosts p to 8 approvals, passing r_{10}'s 7 approvals. But in the process, r_3 rode the wave and was boosted to 9, and so p does not beat r_3. That is, the argument of the previous paragraph worked because the candidates we argued we could ignore were strictly less approved, currently, than the "dangerous" voter we left in that outflanked them. Just to be utterly specific, the first dangerous candidate to the right of p will be the leftmost candidate to the right of p that from V gets at least as many approvals as p (if we were in the nonunique-winner model, we would have said "strictly more approvals than p"). In the figure, that is r_3. And the next dangerous candidate is whichever candidate is the next candidate to the right of that candidate that from V has strictly more approvals than that candidate. In our figure, that is r_6. And so on. And analogously on the left.

So, now that we've set out who the dangerous rivals are, let us see if there is a choice of control action that can let us beat them. In this proof, we'll show just how to beat the very first of them, and then will mention how to continue the process to (attempt to) defeat the rest of them. Crucially, the path we'll take to defeat our first rival will be carefully chosen so that it is the safest possible path—that is, if there is any overall successful control action, then there will be one consistent with the actions we take to beat the first rival (and the iterative repeating of the scheme we outline in this proof will in fact then find such actions).

So let us focus on the first dangerous rival. In our figure, that would be r_3. (Note: If there are no dangerous rivals, then p is already a unique winner, and we're done. If the only dangerous rivals we have are on the left, then mirror image the universe, i.e., reverse L, and now all our dangerous rivals are on the right.) Now, to become the unique winner, p must strictly beat r_3. But note that among all the votes in V', the only ones that can possibly make that happen are (keeping in mind that we consider V' to now have in it only votes that approve of p) *those votes in V' whose right endpoint falls in the half-open interval* $[p, r_3)$. Call this set B. These are the only votes in V' that let p gain approvals relative to r_3. So, in any control action at all that causes p to strictly beat (and we mention in passing that if we were in the nonunique-winner model, we'd here say "tie or beat") r_3, we must put enough votes from that collection, B, in to make that happen. So our only question is which of these votes to add (if there even are enough of them to add to lead to success). We do so as follows. We'll add votes from among these *starting first with the rightmost left endpoint*. This is a perfectly safe strategy, as that choice helps the same or strictly fewer "left" dangerous candidates than would choosing a vote with a left endpoint that goes more to the left. As to the worry that by within B looking only at the left endpoint we may favor some *right-side* candidates that we might have otherwise slighted, that is so, but it is not a worry. We indeed are treating all endpoints between p and just before r_3 as a big blurred class, but the differences among those endpoints affect just the (necessarily nondangerous) candidates in-between, namely, in our figure, candidates r_1 and r_2, but we already argued that by handling r_3 we'll handle both of those, so giving extra approvals to one or both of those (that we could have avoided by making another choice within B) is not an issue. In short we've intentionally smeared our right-end finickiness because it doesn't matter, but

12

left-end details matter a lot. So, in our figure, we need p to gain two approvals relative to r_3. Using the scheme/ordering just mentioned, we'll among the votes in B keep eating through votes until we either achieve this or until we've hit the control problem's limit on the total number of votes we are allowed to add from V'. In this case, in our figure's example, the votes belonging to B (let us assume that the "etc." members of V' contain no members of B) are those votes contained inside the types $[l_4, p]$, $[l_3, r_1]$, and $[l_1, r_2]$. The first vote type to draw on among these is (not withstanding that doing so helps r_2, unlike the case for the other two members of B) $[l_1, r_2]$, as it has the rightmost left endpoint. That vote type is of multiplicity 1, so we take that one vote. The next vote type to draw on is $[l_3, r_1]$. Even though that vote type is of multiplicity 2, p only needs one more approval to beat r_3, so we add to our election just one of the two of the votes of type $[l_3, r_1]$.

And so we have defeated our first dangerous rival, and in the safest possible way—a way that preserves a path to success if any exists.

The completion of the algorithm is basically to do this iteratively as long as candidates remain that (relative to the current state of the votes) are dangerous. So, after we remove a rival, we can consider all the added votes as if they are part of V now, and can redraw the figure, and can move forward, and do all the above against the next dangerous candidate on the right (of course, decreasing the limit on how many votes we are allowed to add by the number we just added), or if we run out of candidates on the right, we can mirror image L and start eating through any remaining dangerous candidates. (Regarding who the dangerous candidates are, the recomputation just mentioned is natural. But we mention in passing that the set of right-side dangerous candidates we'll deal with in the algorithm is in fact the same as the set initially identified—the picture updating won't change it. However, while handling those candidates, it is possible that some of the additional approvals given to p will help it defeat some of the original left-side rivals, as a free side effect.) If we achieve our control goal before burning through the k allowed additional votes from V' (recall that k is part of the problem's input) then we have achieved our control goal, and otherwise it holds that our control goal cannot be achieved.

Let us now turn to the case of deleting voters. Since this case can be handled essentially symmetrically to the adding-voters case (except that we now have a very different choice of which voters to focus on and how to sort them in order to decide which to delete), we merely highlight the differences between the two cases.

In the deleting-voters case, a picture showing our initial state would show just the bar chart of Figure 2, i.e., the number of approvals from V each candidate gets initially. After deletion, the deleted votes would show up below the bar chart, each again as an interval corresponding to the candidates approved of by this voter. However, since it would be pointless to delete a voter approving of p, all such intervals do not include p.

Now, we try to make p strictly beat the next dangerous rival to the right, starting with r_3 in the example shown in Figure 2 (ignoring, of course, the intervals indicating the votes in V' and their multiplicities, as there is no V' in the deleting-voters case). To do so, we now consider all the (remaining, i.e., not yet deleted in the current iteration of our algorithm) votes with intervals whose left endpoint falls in the half-open interval $(p, r_3]$ (in the first

13

iteration in our example). Now we sort these votes *from the rightmost right endpoint to the leftmost right endpoint.* As in the adding-voters case (except there we added votes), we now delete (following the sorted order) just enough votes so that p has one more approval than r_3. If that is impossible, we cannot control. But if it is possible, we've succeeded in making p defeat r_3 in the safest possible way. Recomputing the figure (where, in the redrawn picture, updating the number of the candidates' approvals implies that r_3 no longer is a dangerous candidate for p), we try to keep defeating p's remaining dangerous rivals, and when we run out of dangerous candidates to the right of p, we again mirror image the universe by reversing L and try to make p defeat any remaining dangerous candidates. Eventually, this process ends when we either succeed in making p the unique winner by defeating all of p's dangerous rivals, or we exceed the deletion limit k before reaching this goal. All details not discussed here are obvious from the adding-voters proof.   ❑

We turn from approval voting to plurality voting. Plurality voting's constructive control complexity was studied in detail in the seminal control paper of Bartholdi, Tovey, and Trick [BTT92], and the destructive cases were added by [HHR07] (see [FHHR09a] for the limited adding of candidates cases). For plurality, the situation is close to backwards from approval. Regarding all the adding/deleting control cases defined in this paper, the four voter cases are vulnerable but all six candidate cases are resistant. However, our results here again are in keeping with our paper's theme: All six of these cases become vulnerable for single-peaked societies.

Theorem 3.2 *For the single-peaked case, plurality voting is vulnerable to constructive and destructive control by adding candidates, by adding unlimited candidates, and by deleting candidates, and these results hold in both the unique-winner model and the nonunique-winner model.*[5]

Our proofs of the different parts of this theorem vary greatly in approach. One approach that is particularly useful is dynamic programming. We defer the proof of Theorem 3.2 to Section 3.2.

For two very important voting systems—plurality and approval—we have seen that in every single adding/deleting case where they are known to have NP-hardness complexity shields, the complexity shield evaporates for societies with single-peaked preferences. In the coming manipulation section we will also see a number of cases where NP-hardness shields melt away for single-peaked societies, but we will also see some cases—and earlier such a case was found by Walsh [Wal07]—where existing NP-hardness shields remain in place even if one adds the restriction to a single-peaked society.

The previous paragraph raises the following natural question: Given that restricting to single-peaked preferences can sometimes remove complexity shields, e.g., Theorem 3.1,

---

[5]This result holds in our settled model in which the linear order specifying the society's order on the candidates is part of the input, and also holds in the model in which the linear order is not part of the input but rather the question is whether there exists any valid linear order relative to which there is a way of achieving our control goal.

and can sometimes leave them in place, e.g., Theorem 4.3, can restricting to single-peaked preferences ever erect a complexity shield? It would be very tempting to hastily state that such behavior is impossible. After all, the single-peaked case if anything thins out the flood of possible control-actions/manipulations. But we're talking about complexity here, and fewer options doesn't always mean a less complex problem. In particular, one of those manipulation/control-action options that single-peakedness takes off the table might have been a single, sure-fire path to victory under some election system. In fact, using precisely that approach, and a few other tricks, we have built an artificial election system for which unweighted constructive manipulation by size-3 coalitions is in P in the general case but is NP-complete for the single-peaked case. The system's votes are approval vectors. The system is highly unnatural, highly unsatisfying, and would never be considered for real-world use. However, its goal is just to show that (in a stomach-turningly contrived way) restricting to single-peakedness in concept can, perhaps surprisingly, raise complexity. We conjecture that that strange behavior will never be seen for any existing, natural election system. We summarize this paragraph in the following (manipulation) theorem.

Theorem 3.3 *There exists an election system $\mathcal{E}$, whose votes are approval vectors, for which constructive size-3-coalition unweighted manipulation is in* P *for the general case but is* NP-*complete in the single-peaked model.*[6]

Proof. Let $C = \{c_1, \ldots, c_m\}$ be a set of candidates that is lexicographically ordered (i.e., c_1 is the smallest candidate in this ordering; note that this lexicographical ordering is not related to the society's ordering in single-peaked electorates) and let V be a collection of approval vectors over C. Given C and V, election system $\mathcal{E}$ works as follows. If $\|C\| \leq 2$, all candidates win. Otherwise (i.e., if $\|C\| \geq 3$) we have two possibilities:

1. If there exists no permutation making the given votes single-peaked (which, by Theorem 2.1, can be tested in polynomial time) then all candidates win.[7]

2. Otherwise (i.e., if there is some permutation witnessing single-peakedness of the votes),[8] if the lexicographically smallest candidate, c_1, when viewed as encoding a

---

[6]This result holds both in our settled model where L, the society's ordering of the candidates, is part of the input and in the model where L is not given and the question is whether there is any societal ordering of the candidates that allows the manipulators to succeed.

[7] This feature of the system allows a manipulating coalition to in the general (single-peakedness not required) case cast votes precluding single-peakedness, an attack the single-peaked case does not allow.

[8]For the single-peaked case, our model also allows the election system $\mathcal{E}$ to directly access the input permutation relative to which V is single-peaked. For "real" election systems this would be irrelevant but regarding building an artificial counterexample system, such as $\mathcal{E}$, this may be important, although Theorem 3.3 holds in either case. Assuming that the election system knows the society's permutation witnessing single-peakedness, as our model does, matches the intuition that the societal ordering of candidates is open information. Also, this assumption is fairer considering that our manipulation algorithms do get direct access to the input permutation of society, and our model is that that permutation is fixed (in candidate/voter addition cases it already must handle *all* candidates and *all* voters, and in candidate/voter deletion cases it must stick with the permutation even if after deletion other permutations are consistent with the remaining candidates/voters). The permutation reflects "society" and does not change from what the input says it is.

15

boolean formula encodes a satisfying one and the lexicographically smallest approval vector is all zeros, then all candidates win, else all candidates lose.

For the general case, the constructive size-3-coalition unweighted manipulation problem is in P: Always have the three manipulators cast the following votes (with respect to the lexicographic ordering of $C = \{c_1, \ldots, c_m\}$):

$$(1, 1, 0, 0, \ldots, 0)$$
$$(0, 1, 1, 0, \ldots, 0)$$
$$(1, 0, 1, 0, \ldots, 0)$$

ensuring that the electorate is not single-peaked, and so all candidates win, including the distinguished candidate.

For the single-peaked case, however, the constructive size-3-coalition unweighted manipulation problem is NP-complete. The problem is clearly in NP, and to show NP-hardness we reduce the well-known NP-complete boolean formula satisfiability problem to it. Given a boolean formula φ (encoded as a string in $\{0, 1\}^*$), map it to the instance (C, V, φ) of the constructive size-3-coalition unweighted manipulation problem, where $C = \{\varphi, 1^{|\varphi|+1}, 1^{|\varphi|+2}\}$, V consists of one nonmanipulator who approves of all candidates and three manipulators, and where φ is the distinguished candidate the manipulators want to make win. We claim that they can succeed in doing so exactly if φ is satisfiable. Indeed, if φ is satisfiable then φ can be made a winner by one of the manipulators casting the vote $(0, 0, 0)$. On the other hand, if φ is not satisfiable (and the electorate is restricted to being single-peaked) then φ cannot win, no matter which votes the manipulators cast. ❑

3.2   The Deferred Proof of Theorem 3.2

In this section we provide, as a series of lemmas that cover appropriate cases, our proof of Theorem 3.2. One of the reasons why candidate-control problems for plurality are hard in the general case (i.e., when the voters are not restricted to be single-peaked) is that then adding or deleting even a single candidate can possibly affect the scores of all the other candidates. However, in the single-peaked case adding or deleting a candidate can affect at most two other candidates. This observation, formalized in the next lemma, is at the heart of the vulnerability proofs of this section.

Lemma 3.4 *Let (C, V) be an election where $C = \{c_1, \ldots, c_m\}$ is a set of candidates, V is a collection of voters whose preferences are single-peaked with respect to a linear order L, and where $c_1 \, L \, c_2 \, L \, \cdots \, L \, c_m$. Within plurality, if $m \geq 2$ then*

1. *$score_{(C,V)}(c_1) = score_{(\{c_1,c_2\},V)}(c_1)$,*

2. *for each i, $2 \leq i \leq m-1$, $score_{(C,V)}(c_i) = score_{(\{c_{i-1},c_i,c_{i+1}\},V)}(c_i)$, and*

3. *$score_{(C,V)}(c_m) = score_{(\{c_{m-1},c_m\},V)}(c_m)$.*

Proof. Let (C,V), L, and m be as in the statement of the lemma. We prove the second case, $2 \leq i \leq m-1$ (the proofs of the boundary cases are analogous). Fix an integer i, $2 \leq i \leq m-1$. Each voter v in C that ranks c_i as her top choice, clearly, still prefers c_i to both c_{i-1} and c_{i+1} when limited to the choice between c_i, c_{i-1}, and c_{i+1}. On the other hand, consider a voter v in V who ranks some candidate c_j, $c_i \neq c_j$, as her most preferred candidate. If $j \leq i-1$ then, due to the fact that preferences are single-peaked (or, $j = i-1$), this voter prefers c_{i-1} to c_i. If $j \geq i+1$ then, for the same reason, this voter prefers c_{i+1} to c_i. In either case, this voter does not give her point to c_i. As a result, $score_{(C,V)}(c_i) = score_{(\{c_{i-1},c_i,c_{i+1}\},V)}(c_i)$. ❏

Using Lemma 3.4, we show that in the single-peaked case plurality is vulnerable to constructive control by adding candidates.

Lemma 3.5 *For the single-peaked case, plurality voting is vulnerable to constructive control by adding candidates both in the unique-winner model and in the nonunique-winner model.*

Proof. We give a polynomial-time algorithm for the single-peaked variant of the constructive control by adding candidates problem. We focus on the unique-winner case, but the proof can easily be adapted to the nonunique-winner model and we will be pinpointing the necessary changes throughout the proof. Our input is an election $(C \cup A, V)$—where $C = \{p, c_1, \ldots, c_{m'}\}$ and $A = \{a_1, \ldots, a_{m''}\}$—a candidate p in C, a nonnegative integer k, and a linear order L over $C \cup A$.[9] After checking that the voters in V are single-peaked with respect to L, we ask if it is possible to find a subset A' of A such that (a) $\|A'\| \leq k$ and (b) for each $c \in C \cup A'$, $c \neq p$, it holds that $score_{(C \cup A',V)}(c) < score_{(C \cup A',V)}(p)$ $(score_{(C \cup A',V)}(c) \leq score_{(C \cup A',V)}(p)$, in the nonunique-winner case). Let $D = C \cup A$ and let us rename the candidates so that $D = \{d_1, \ldots, d_m\}$, where $m = 1 + m' + m''$, and $d_1 L d_2 L \cdots L d_m$. Without loss of generality we can assume that p is neither d_1 nor d_m. If this were not the case, we could add to C two dummy candidates, d' and d'', that each voter ranks as the least desirable ones and such that $d' L d_1 L \cdots L d_m L d''$.

Before we describe the algorithm formally, let us explain it intuitively. Let us assume that our instance is a "yes" instance and that A' is a subset of A witnessing this fact. There must be two candidates in $C \cup A'$, d_ℓ and d_r, that are direct neighbors of p with respect to L restricted to $C \cup A'$ (i.e., $d_\ell L p L d_r$ and there is no $d_i \in C \cup A'$ such that $d_i \neq p$ and $d_\ell L d_i L d_r$). By Lemma 3.4, the score of p is "fixed" by d_ℓ and d_r in the sense that $score_{(C \cup A',V)}(p) = score_{(\{d_\ell,p,d_r\},V)}(p)$. Also, given such d_ℓ and d_r we can view $C \cup A'$ as partitioned into the left part, $C_\ell \cup A'_\ell$, containing d_ℓ and the candidates preceding her, the

---

[9]If we are in the model in which no linear societal order is given simply use, for example, the algorithm of Escoffier, Lang, and Öztürk [ELÖ08] to find one relative to which our input is single-peaked. If none exists, state that and terminate. If the algorithm of [ELÖ08] gives us some valid order, it is not hard to see that the algorithm we are about to describe in this proof will find a successful control action exactly if there is any valid societal linear order relative to which a successful control action exists. (In the remaining proofs of Section 3.2, we'll handle only the case that the societal order witnessing single-peakedness is given in our polynomial-time algorithms showing vulnerability claims. Note, however, that this footnote's argument handles the case of such a societal order not being given analogously in each such case.)

right part, $C_r \cup A'_r$, containing d_r and the candidates succeeding her, and the middle part, containing p. The candidates in the left part and the candidates in the right part all have scores lower than p.

The idea of the algorithm is to try all possible pairs of candidates d_ℓ and d_r that can be direct neighbors of p and to test, for each such pair, if it is possible to add candidates so that (a) each of the candidates in the left part has a score lower than p, (b) each of the candidates in the right part has a score lower than p, and (c) the number of candidates added is at most k. This approach is formalized in the algorithm below.

1. For each pair of candidates, d_ℓ and d_r, such that

    (a) $d_\ell \, L \, p \, L \, d_r$, and

    (b) there is no candidate $c_i \in C$ such that $d_\ell \, L \, c_i \, L \, d_r$,

    execute Steps 2 through 5 below. After trying all pairs d_ℓ, d_r without accepting, reject.

2. (a) Set $b' = score_{(\{d_\ell, p, d_r\}, V)}(p) - 1$.

    (b) Set $C_\ell = \{c_i \mid c_i \, L \, d_\ell\} \cup \{d_\ell\}$ and $C_r = \{c_i \mid d_r \, L \, c_i\} \cup \{d_r\}$.[10]

    (c) Set $A_\ell = \{a_i \mid a_i \, L \, d_\ell\}$ and $A_r = \{a_i \mid d_r \, L \, a_i\}$.

    (d) V_ℓ is the collection of those voters in V who, among the candidates in $C_\ell \cup C_r \cup A_\ell \cup A_r \cup \{p\}$, rank first either d_ℓ or a candidate d_i such that $d_i \, L \, d_\ell$; V_r is the collection of those voters in V who, among the candidates in $C_\ell \cup C_r \cup A_\ell \cup A_r \cup \{p\}$, rank first either d_r or a candidate d_i such that $d_r \, L \, d_i$.[11]

3. Find a smallest (with respect to cardinality) set B_ℓ, $B_\ell \subseteq A_\ell$, such that for each $c \in C_\ell \cup B_\ell$ it holds that $score_{(C_\ell \cup B_\ell, V_\ell)}(c) \le b'$. If such a set B_ℓ does not exist then discard the current pair (d_ℓ, d_r) and try the next one.

4. Find a smallest (with respect to cardinality) set B_r, $B_r \subseteq A_r$, such that for each $c \in C_r \cup B_r$ it holds that $score_{(C_r \cup B_r, V_r)}(c) \le b'$. If such a set B_r does not exist then discard the current pair (d_ℓ, d_r) and try the next one.

5. If $\|B_\ell\| + \|B_r\| + \|A \cap \{d_\ell, d_r\}\| \le k$ then accept (the set A' that we seek is $B_\ell \cup B_r \cup (\{d_\ell, d_r\} \cap A)$).

It is easy to see that the above algorithm is correct and that it also works in the nonunique-winner model if we replace Step 2a with $b' = score_{(\{d_\ell, p, d_r\}, V)}(p)$. We will now show that the algorithm runs in polynomial time. There are at most quadratically many (with respect to $m = \|D\|$) pairs (d_ℓ, d_r) to try, so each of the above steps will be executed

---

[10]Yes, we put d_ℓ in C_ℓ and d_r in C_r even if either of them belongs to A.

[11]Note that if B_ℓ is a subset of A_ℓ and B_r is a subset of A_r then in election $(C \cup B_\ell \cup B_r \cup \{d_\ell, d_r\}, V)$ it holds that (a) each voter in V_ℓ ranks first some candidate in $C_\ell \cup B_\ell \cup \{d_\ell\}$, (b) each voter in V_r ranks first some candidate in $C_r \cup B_r \cup \{d_r\}$, and (c) all the remaining voters rank p first.

at most polynomially often. It remains to show that Steps 3 and 4 can each be executed in polynomial time. To this end, we define the following functional problem.

Definition 3.6 *In the DemoteByAddingCandidates problem we are given an election $(C \cup A, V)$, where the voters in V are single-peaked with respect to a given order L, and a nonnegative integer b. We ask what the size is of a smallest set $A' \subseteq A$ such that for each candidate $c \in C \cup A'$ it holds that $score_{(C \cup A', V)}(c) \leq b$. If such a set A' does not exist then, by convention, the answer is ∞.*

Computing B_ℓ (B_r) in Step 3 (in Step 4) simply requires solving an instance of DemoteByAddingCandidates with election $(C_\ell \cup A_\ell, V_\ell)$ (election $(C_r \cup A_r, V_r)$), order L (restricted to an appropriate subset of candidates), and integer $b = b'$. Fortunately, DemoteByAdding-Candidates is easy to solve.

Lemma 3.7 *DemoteByAddingCandidates is computable in polynomial time.*

Proof. We give a polynomial-time algorithm for the DemoteByAddingCandidates problem. Our input consists of an election $(C \cup A, V)$, a nonnegative integer b, and a linear order L over $C \cup A$, where $C = \{c_1, \ldots, c_{m'}\}$, $A = \{a_1, \ldots, a_{m''}\}$, and voters in V have single-peaked preferences with respect to the order L.[12]

Let $D = C \cup A$ and let us rename the candidates so that $D = \{d_1, \ldots, d_m\}$, where $m = m' + m''$ and $d_1\, L\, d_2\, L \cdots L\, d_m$. Without loss of generality, we can assume that $d_1 \in C$ and that $m \geq 3$.[13] Let i, j, and k be three positive integers such that $1 \leq i < j < k \leq m$. We define

1. $s(i, j, k) = score_{(\{d_i, d_j, d_k\}, V)}(d_j)$,

2. $s(0, j, k) = score_{(\{d_j, d_k\}, V)}(d_j)$, and

3. $s(i, j, m+1) = score_{(\{d_i, d_j\}, V)}(d_j)$.

Note that, by Lemma 3.4, if B is a subset of A such that (a) $\{d_i, d_j, d_k\} \subseteq C \cup B$ and (b) there is no candidate d_l ($l \neq j$) such that $d_i\, L\, d_l\, L\, d_k$, then $s(i, j, k) = score_{(C \cup B, V)}(d_j)$.

Let i and j be two positive integers such that $1 \leq i < j \leq m$. By $A(i, j)$ we mean a family of subsets of A such that a set B, $B \subseteq A$, belongs to $A(i, j)$ if and only if (a) $d_i \in C \cup B$, (b) $d_j \in C \cup B$, and (c) $(C \cup B) \cap \{d_{i+1}, \ldots, d_{j-1}\} = \emptyset$.

We now define a key value for our algorithm. Let i and j be two positive integers such that $1 \leq i < j \leq m$. We define

$$f(i, j) = \min_{B \in A(i,j)} \left\{ \|B\| \mid (\forall d_\ell \in C \cup B, 1 \leq \ell \leq i)[score_{(C \cup B, V)}(c_\ell) \leq b] \right\}.$$

---

[12]Keep in mind that within the proof of this lemma, C, A, V, L, and b refer to the input to the DemoteByAddingCandidates problem, not to the input to the Control by Adding Candidates problem.

[13]If m were less than 3, we could add to C an appropriate number of dummy candidates, each ranked by all the voters as less desirable than any of the candidates in $C \cup A$. Extending L so that for each added dummy candidate d it holds that $d\, L\, d_1$ and the voters are single-peaked with respect to this extended L is easy.

We adopt the convention that $\min \emptyset$ is ∞. It is easy to verify that the algorithm's output should be

$$\min_{1 \leq i < j \leq m} \{f(i,j) \mid \{d_{j+1}, \ldots, d_m\} \subseteq A \wedge s(i,j,m+1) \leq b\}.$$

It remains to show that we can compute the values $f(i,j)$, $1 \leq i < j \leq m$, in polynomial time. The rest of the proof is devoted to this task.

We compute the values $f(i,j)$ using dynamic programming. Let j be a positive integer $1 < j \leq m$. It is easy to see that

$$f(1,j) = \begin{cases} \infty & \text{if there is some } d_i \in C \text{ such that } 1 < i < j, \\ \infty & \text{if the above does not hold but } s(0,1,j) > b, \\ \|A \cap \{d_1, d_j\}\| & \text{otherwise.} \end{cases}$$

Given these initial values,[14] we can compute arbitrary values $f(j,k)$, $1 < j < k \leq m$, using the following facts. Let us fix j and k such that $1 < j < k \leq m$. If there is a $d_\ell \in C$ such that $j < \ell < k$ then $f(j,k) = \infty$ (because $A(j,k)$ is empty). Otherwise

$$f(j,k) = \min_{1 \leq i < j} \{f(i,j) + \chi_A(d_k) \mid s(i,j,k) \leq b\},$$

where χ_A is the characteristic function of A (i.e., $\chi_A(d_k)$ is 1 if $d_k \in A$, and is 0 otherwise). It is easy to see that using standard dynamic-programming techniques and the above equations we can compute $f(i,j)$ for arbitrary positive integers i,j, $1 \leq i < j \leq m$, in polynomial time. ❑ Lemma 3.7

This completes the proof that for the case of single-peaked voters, plurality is vulnerable to constructive control by adding candidates. ❑ Lemma 3.5

As a corollary to Lemma 3.5, we obtain the corresponding result for the case of adding an unlimited number of candidates, because for each unlimited case it is sufficient to use the algorithm for the corresponding limited case with the limit on the number of candidates that can be added set to equal the number of spoiler candidates.

Corollary 3.8 *For the single-peaked case, plurality voting is vulnerable to constructive control by adding an unlimited number of candidates in both the unique-winner model and the nonunique-winner model.*

In the following proofs (both in the destructive adding-candidates cases and in the deleting-candidates cases) in addition to Lemma 3.4 we will need one more tool, namely, the notion of a neighborhood of a candidate c in an order L witnessing single-peakedness of an electorate.

Definition 3.9 *Let (C,V) be an election and let L be an order such that the voters in V are single-peaked with respect to L. Let c be a candidate in C. We*

---

[14]Note that here it is important that $d_1 \in C$. If this were not the case, we would also have to directly compute the values $f(i,j)$ for some i,j where $1 < i < j$.

rename the candidates in C such that $C = \{b_{m'}, \ldots, b_2, b_1, c, d_1, d_2, \ldots, d_{m''}\}$ and $b_{m'} L \cdots L b_2 L b_1 L c L d_1 L d_2 L \cdots L d_{m''}$. For each two positive integers i and j we set $D_{ij}(C, c) = \{b_1, \ldots, b_{\min(i,m')}\} \cup \{d_1, \ldots, d_{\min(j,m'')}\}$. We define the direct neighborhood of a candidate c to be the set $D(C, c) = \{D_{ij}(C, c) \mid 0 \le i \le m' \text{ and } 0 \le j \le m''\}$.[15]

Clearly, given a single-peaked election (C, V) and a candidate $c \in C$, $\|D(C, c)\|$ is polynomial in $\|C\|$. We now turn to the destructive adding-candidates cases in both the unique-winner model and the nonunique-winner model. The proof is much simpler than in the constructive cases.

Lemma 3.10 *For the single-peaked case, plurality voting is vulnerable to destructive control by adding candidates in both the unique-winner model and the nonunique-winner model.*

Proof. Let $(C \cup A, V)$ be an election where $C = \{d, c_1, \ldots, c_m\}$, $A = \{a_1, \ldots, a_{m'}\}$, and V is a collection of voters whose preference orders are single-peaked with respect to a given order L. The candidates from C are already registered and the candidates from A can be added (i.e., A is the spoiler candidate set). Let k be some nonnegative integer. We now give a polynomial-time algorithm that tests whether there exists a set $A' \subseteq A$ such that (a) $\|A'\| \le k$ and (b) d is not a unique winner of $(C \cup A', V)$.

The algorithm works as follows: For each up-to-3-element subset A'' of A, check if d is a unique winner of $(C \cup A'', V)$ and accept if this is not the case (for the first such A'' found). If we have not accepted for any such subset A'' then reject. This algorithm clearly works in polynomial time. It remains to show that it is correct. If for each $A' \subseteq A$ it holds that d is a unique winner of $(C \cup A', V)$ then the above algorithm (correctly) rejects. On the other hand, let us assume that for some $A' \subseteq A$ it holds that d is not a unique winner of $(C \cup A', V)$ and let us fix one such A'. We claim that there is an up-to-size-3 subset A'' of A such that d is not a unique winner of $(C \cup A'', V)$.

Let $B = D_{1,1}(C \cup A', d)$, that is, B contains the direct neighbors of d (among the candidates in $C \cup A'$, with respect to L). By definition, $\|B\| \le 2$, and, from Lemma 3.4, $score_{(C \cup A', V)}(d) = score_{(B \cup \{d\}, V)}(d)$. Since d is not a unique winner of $(C \cup A', V)$, there is some candidate $c \in C \cup A'$, $c \ne d$, such that $score_{(C \cup A', V)}(c) \ge score_{(C \cup A', V)}(d)$. Define $A'' = (B \cup \{c\}) \cap A$. Clearly, $\|A''\| \le 3$ and, since B is a "radius-1" neighborhood of d, we have $score_{(C \cup A'', V)}(d) = score_{(B \cup \{d\}, V)}(d)$. On the other hand, since $C \cup A'' \subseteq C \cup A'$ and c belongs to both sets, it holds that $score_{(C \cup A'', V)}(c) \ge score_{(C \cup A', V)}(c)$. It follows that in election $(C \cup A'', V)$ candidate c has a score at least as high as that of d and, thus, d is not a unique winner of $(C \cup A'', V)$.

It is easy to see that the same approach can be used to solve the nonunique-winner case. Our algorithm now tries all up-to-3-element subsets A'' of A and accepts if and only if there is one such that d is not a winner of $(C \cup A'', V)$. The proof of correctness is analogous to the unique-winner case. This completes the proof. ❑

---

[15]The notation of this definition, in essence, ties the order L to the candidate set C. We will sometimes wish to speak of, e.g., $D_{ij}(C', c)$ for some $C' \subseteq C$, $c \in C'$, and this notation allows us to naturally, implicitly speak of an order L' induced by C'.

As an easy corollary we obtain that in the single-peaked case plurality is also vulnerable to destructive control by adding unlimited candidates.

Corollary 3.11 *For the single-peaked case, plurality voting is vulnerable to destructive control by adding an unlimited number of candidates in both the unique-winner model and the nonunique-winner model.*

We now turn to the constructive and destructive deleting-candidates cases in both the unique-winner model and the nonunique-winner model. The main idea of our algorithms for the deleting-candidates cases is that it is sufficient to focus on deleting candidates adjacent to the preferred one (to increase her score) and then to delete those remaining candidates that still defeat the preferred one.

Lemma 3.12 *For the single-peaked case, plurality voting is vulnerable to constructive and destructive control by deleting candidates in both the unique-winner model and the nonunique-winner model.*

Proof. We give a polynomial-time algorithm for the single-peaked case of constructive control by deleting candidates for plurality. We focus on the unique-winner model but it is easy to see how to modify our algorithm to work in the nonunique-winner model.

Our algorithm's input contains an election (C, V), a preferred candidate $p \in C$, a nonnegative integer k (the number of candidates that we are allowed to delete), and an order L such that the voters in V are single-peaked with respect to L.

By Lemma 3.4, we see that the only possible scores that p can have after deleting some candidates are in the set $S = \{score_{(C-D,V)}(p) \mid D \in D(C, p)\}$ and that each of these scores can be obtained by deleting some subset $D \in D(C, p)$ of the candidates. In consequence, to ensure that p is the unique winner we first need to delete some subset $D \in D(C, p)$ of candidates and then—since it is impossible to decrease a candidate's score via deleting other candidates—delete all those candidates that still have more points than p. Of course, we do not know *a priori* which set D in $D(C, p)$ to delete, but since $\|D(C, p)\|$ is polynomially bounded in $\|C\|$, we can try all its members.

Formally, our algorithm works as follows.

1. For each $D \in D(C, p)$, execute Step 2. If, after trying all D's, we do not accept, reject.

2. (a) While there exists a candidate $c \in C - D$, $c \neq p$, such that $score_{(C-D,V)}(c) \geq score_{(C-D,V)}(p)$, add c to D.
   (b) If $\|D\| \leq k$ then accept.

Clearly, this algorithm works in polynomial time and, by the preceding discussion, it accepts if and only if it is possible to ensure p's victory via control by deleting candidates. For the nonunique-winner case, we simply need to change the inequality in Step 2a from "$\geq$" to "$>$."

In the destructive case we are given an election (C, V), a despised candidate $d \in C$, a nonnegative integer k (the number of candidates we can delete), and an order L such that the voters in V are single-peaked with respect to L. (We assume that d is a unique plurality winner of (C, V); otherwise, we can immediately accept.) It is sufficient to find a candidate c and a subset $D \subseteq C - \{d\}$ of candidates such that $\|D\| \leq k$ and $score_{(C-D,V)}(c) \geq score_{(C-D,V)}(d)$. As in the constructive case, by Lemma 3.4, it suffices to try each candidate $c \in C - \{d\}$ and each set $D \in D(C, c)$. That is, if there is a candidate $c \in C - \{d\}$ and a set $D \in D(C, c)$ such that (a) $\|D\| \leq k$, (b) $d \notin D$, and (c) $score_{(C-D,V)}(c) \geq score_{(C-D,V)}(d)$, then we accept. Otherwise we reject. Clearly, this algorithm works in polynomial time and, by Lemma 3.4 and the fact that deleting candidates does not decrease the remaining candidates' scores, is correct. The algorithm can be modified in an obvious way to work for the nonunique-winner model. ❑

4 Manipulation

In this section we study constructive coalition weighted manipulation. Recall that in the single-peaked case the manipulators must cast votes that are consistent with the linear ordering of the candidates (which is part of the input in the single-peaked case) that defines the society's single-peakedness. However, all our "single-peaked case is in P" results in this section also hold in a model in which the linear order of the candidates is not part of the input and we will comment on these cases in our proofs below.

This paper's theme is that single-peakedness removes many NP-hardness shields, and for manipulation we show that via Theorem 4.1 (and its implications within Theorem 4.2) and Theorem 4.4. Note that in the general case, the election systems of parts 1 and 3 of Theorem 4.1 and the "$k_1 \geq 2 \wedge k_0 \geq 1$" cases of part 2 of this theorem are known to be NP-complete [HH07,PR07,CSL07], and the remaining part 2 cases are easily seen to be in P [HH07].)

Theorem 4.1 *For the single-peaked case, the constructive coalition weighted manipulation problem (in both the nonunique-winner model and the unique-winner model) for each of the following election systems is in* P*:*

1. *The scoring protocol $\alpha = (2, 1, 0)$, i.e., 3-candidate Borda elections. (This is a special case of Theorem 4.4.)*

2. *Each scoring protocol $\alpha = (\overbrace{1, \ldots, 1}^{k_1}, \overbrace{0, \ldots, 0}^{k_0})$, $k_1 \geq k_0$. (This includes a variety of ℓ-veto and ℓ'-approval protocols, e.g., the 3-veto cases for $m \geq 6$ candidates in Theorem 4.2.)*

3. *Veto.[16]*

---

[16]All three of Theorem 4.1's cases hold both in our settled model where L, the society's ordering of the candidates, is part of the input and in the model where L is not given and the question is whether there is any societal ordering of the candidates that allows the manipulators to succeed.

Proof. For each of the election systems stated in the theorem, we give a polynomial-time algorithm solving the constructive coalition weighted manipulation problem for the single-peaked case. Each of these algorithms takes as input a list C of candidates, a list S of nonmanipulative voters each specified by a single-peaked preference over C (with respect to the given linear ordering L on C that defines single-peakedness) and an integer weight, a list of the weights of the voters in our manipulating coalition T, and the candidate $p \in C$ that the manipulators in T seek to make a winner (a unique winner). Note that all manipulators in T must also have single-peaked preferences with respect to L.

The scoring protocol $\alpha = (2, 1, 0)$. We start with the nonunique-winner case. Let $C = \{a, b, p\}$. Without loss of generality, it is enough to consider the following two societal orderings L. (If L is not part of the input, we simply try all societal orders.)

Case 1: $a \, L \, p \, L \, b$. Due to single-peakedness, only the following votes are allowed with respect to L:

$$p > a > b, \quad p > b > a, \quad a > p > b, \quad b > p > a.$$

Note that, regardless of the weights of the voters in S, we have $score_S(p) \geq score_S(a)$ or $score_S(p) \geq score_S(b)$. So, at most one of a and b beats p, say a. Set all votes in T to $p > b > a$.

Case 2: $p \, L \, a \, L \, b$. In this case, due to single-peakedness, only the following votes are allowed with respect to L:

$$p > a > b, \quad a > p > b, \quad a > b > p, \quad b > a > p.$$

In this case, a best vote for all voters in T is $p > a > b$.

In both cases, we accept if the manipulation action described above was successful (i.e., made p a winner in election $(C, S \cup T)$), and we reject otherwise.

To handle the unique-winner case, note that Case 2 remains the same. For Case 1, even if $score_S(b) = score_S(p)$, setting all voters in T to $p > b > a$ will work and is still optimal. In both cases, for checking success of the manipulation action we now test whether it makes p the unique winner.

Scoring protocols $\alpha = (\overbrace{1, \ldots, 1}^{k_1}, \overbrace{0, \ldots, 0}^{k_0})$, $k_1 \geq k_0$. Let $\alpha = (\overbrace{1, \ldots, 1}^{k_1}, \overbrace{0, \ldots, 0}^{k_0})$, $k_1 \geq k_0$, be a scoring protocol. We will handle the cases $k_1 > k_0$ and $k_1 = k_0$ separately. As above, we will present algorithms for both the nonunique-winner model and the unique-winner model.

First, assume $k_1 > k_0$. In any election with $m = k_0 + k_1$ candidates, if the voters are single-peaked with respect to some linear order L,[17] then the middle candidate(s) (namely, the $\lceil m/2 \rceil$nd candidate in L if m is odd, and the $(m/2)$nd and the $(1 + m/2)$nd candidate in L

---

[17] If L is not part of the input, we simply try all possible societal orderings. This can be done in polynomial time, since there are only a constant number of candidates.

if m is even) will always be among the winners. Thus, given an instance of the constructive coalition weighted manipulation problem, our distinguished candidate p can be made a winner in $(C, S \cup T)$ if and only if p does not lose a point in S. So our algorithm checks whether p does not lose a point in S and accepts or rejects accordingly.

For the unique-winner case within scoring protocol $\alpha = (\overbrace{1, \ldots, 1}^{k_1}, \overbrace{0, \ldots, 0}^{k_0})$, $k_1 > k_0$, p can be made a unique winner if and only if

1. p does not lose a point in S, and

2. for all candidates c that are tied with p in S, it is possible to make c lose a point in T (while keeping T single-peaked with respect to L).

Since the number of candidates is fixed, our algorithm can easily check whether these two conditions hold and then accepts or rejects accordingly.

Now, assume $k_1 = k_0$. Given an instance of the constructive coalition weighted manipulation problem, it is enough to consider a linear order L (which witnesses single-peakedness; also, recall Footnote 17) that ranks at least $k = k_0 = k_1$ candidates before p and fewer than k candidates after p: $c_1 \, L \, c_2 \, L \, \cdots \, L \, c_\ell \, L \, p \, L \, c_{\ell+2} \, L \, \cdots \, L \, c_{2k}$ with $\ell \geq k$. For every candidate c_i, $\ell + 2 \leq i \leq 2k$, to the right of p, we have $score_S(p) \geq score_S(c_i)$, and for every candidate c_j, $k + 1 \leq j \leq \ell$, to the left of p and to the right of the middle, it holds that if p gets a point then so does c_j. It follows that p can be made a winner if and only if p is a winner in $(C, S \cup T)$ when every voter in T ranks the candidates from right to left by $c_{2k} > c_{2k-1} > \cdots > c_{\ell+2} > p > c_\ell > \cdots > c_1$.

For the unique-winner case within scoring protocol $\alpha = (\overbrace{1, \ldots, 1}^{k}, \overbrace{0, \ldots, 0}^{k})$, if p's right neighbor in L is tied with p in S (i.e., $score_S(c_{\ell+2}) = score_S(p)$), then set a voter of lowest weight in T to $p > c_\ell > c_{\ell-1} > \cdots > c_1 > c_{\ell+2} > \cdots > c_{2k}$ and set all other voters in T to rank the candidates from right to left by $c_{2k} > c_{2k-1} > \cdots > c_{\ell+2} > p > c_\ell > \cdots > c_1$. (Note that T is single-peaked with respect to L.) Our algorithm can easily check whether this makes p a unique winner in $(C, S \cup T)$, and if so it accepts, otherwise it rejects.

Veto elections. The proof for veto elections (i.e., for scoring protocols $\alpha = (1, \ldots, 1, 0)$ where the number of candidates varies), similarly as the proof of the previous part of this theorem, hinges on the following observation: Given an instance of the constructive coalition weighted manipulation problem, where the votes are single-peaked with respect to the society's order L, p can be made a winner if and only if p is the only candidate or p is never last in S.

For the unique-winner case, note that since at most two candidates can be ranked last in single-peaked elections, unique winners within veto elections can only exist if $\|C\| \leq 3$. The result for $\|C\| \in \{2, 3\}$ follows from the previous part of the theorem, and the $\|C\| = 1$ case is trivial. This completes the proof of the last part of Theorem 4.1 and the whole proof. ❑

We now come to an unusual case. For general votes (not limited to single-peaked societies), the constructive coalition weighted manipulation problem for 3-veto is in P for three or four candidates, and is NP-complete (and so resistant) for five or more candidates (see [HH07]). (3-veto is not meaningfully defined for two or fewer candidates.) However, for single-peaked votes, 3-veto shows a remarkable behavior: Moving from five to six candidates *lowers* the complexity of this manipulation problem. In some papers, authors state that if one knows the number of candidates (if any) at which a system switches from easy to manipulate to hard to manipulate, it is easy by adding dummy candidates to see that for all larger number of candidates the system remains hard. The following theorem should stand as a caution to take that view only if one has carefully built and checked a "dummy candidates" construction for one's specific case.

Theorem 4.2 *For the single-peaked case, the constructive coalition weighted manipulation problem (in both the unique-winner model and the nonunique-winner model) for m-candidate 3-veto elections is in* P *for $m \in \{3, 4, 6, 7, 8, \ldots\}$ and is resistant (indeed,* NP*-complete) for $m = 5$.*[18]

Proof.   Consider the following four cases.

$\boldsymbol{m = 3.}$ In this case, we are looking at the scoring protocol $(0, 0, 0)$. It is immediate that in this scoring protocol all candidates are always tied for winner, and so p can always be made a winner. For the same reason, p can never be a unique winner.

$\boldsymbol{m = 4.}$ In this case, we are looking at the scoring protocol $(1, 0, 0, 0)$. It is immediate that p can be made a winner (a unique winner) if and only if p is a winner (a unique winner) in the election where every manipulator ranks p first.

$\boldsymbol{m = 5.}$ In this case, we are looking at the scoring protocol $(1, 1, 0, 0, 0)$. It is immediate that the constructive coalition weighted manipulation problem is in NP (simply guess single-peaked votes for the manipulators and verify that p is a winner of the election). To show NP-hardness, we reduce from PARTITION.

Given an instance of PARTITION, i.e., a set $\{k_1, \ldots, k_n\}$ of n distinct positive integers that sum to $2K$, construct the following instance of Constructive Coalition Weighted Manipulation: The set of candidates is $C = \{a, b, c, d, p\}$ with p being our distinguished candidate. The set of nonmanipulators S consists of two voters, each of weight K. One of the voters votes $c > a > p > b > d$ and the other voter votes $d > b > p > a > c$. (Note that these two voters fix the society's order.) We also have a set T of n manipulators. The weights of the manipulators are $k_1, k_2, \ldots, k_n$. We claim that there is a partition if and only if the manipulators can cast single-peaked votes that make p a winner.

---

[18]This result holds both in our settled model where L, the society's ordering of the candidates, is part of the input and in the model where L is not given and the question is whether there is any societal ordering of the candidates that allows the manipulators to succeed.

First suppose that there exists a subset S' of $\{k_1, \ldots, k_n\}$ that sums to K. For every $i \in \{1, \ldots, n\}$, we set the weight k_i manipulator to $p > a > b > c > d$ if k_i in S' and to $p > b > a > c > d$ otherwise. In the resulting election, a, b, and p each score $2K$ points and c and d score K points. It follows that p is a winner of the election.

For the converse, suppose the voters in T vote single-peaked such that p is a winner of the election. Since for every single-peaked vote, if p gets a point then a or b gets a point, it follows that $score_T(p) \leq score_T(a) + score_T(b)$. Since p is a winner of the resulting election, $score_T(p) \geq K + score_T(a)$ and $score_T(p) \geq K + score_T(b)$. It follows that $score_T(p) = 2K$ and that $score_T(a) = score_T(b) = K$. But then the weights of the voters in T that give a point to a sum to K and so we have found a partition.

Changing the weights of the voters in S from K to $K-1$ handles the unique-winner case.

$\boldsymbol{m \geq 6}$. This has already been proven as the special case $k_1 \geq k_0 = 3$ of part 2 of Theorem 4.1.

This completes the proof of Theorem 4.2. ❑

We now present some cases that are known (see [HH07]) to be NP-hard in the general case and that we can prove remain hard even in the single-peaked case.

Theorem 4.3 *For the single-peaked case, the constructive coalition weighted manipulation problem (in both the unique-winner model and the nonunique-winner model) is resistant (indeed,* NP*-complete) for the following scoring protocols:*

1. $\alpha = (3, 1, 0)$.

2. $\alpha = (3, 2, 1, 0)$, i.e., 4-candidate Borda elections.[19]

Proof. We consider the two scoring protocols, $\alpha = (3, 1, 0)$ and $\alpha = (3, 2, 1, 0)$ separately. In each case it is easy to see that the constructive coalition weighted manipulation problem is in NP and it remains to give the NP-hardness proof.

The scoring protocol $\boldsymbol{\alpha = (3, 1, 0)}$. To show NP-hardness, we reduce from PARTITION. Given a set $\{k_1, \ldots, k_n\}$ of n distinct positive integers that sum to $2K$, define the following instance of the constructive coalition weighted manipulation problem: The candidate set is $C = \{a, b, p\}$, there are two nonmanipulative voters in S, one of the form $a > p > b$ and one of the form $b > p > a$ (note that S is single-peaked with respect to $a \, L \, p \, L \, b$, its reverse, and no other order) and both of weight $5K$, and there are n manipulators in T where the ith manipulator has weight k_i, $1 \leq i \leq n$.

---

[19]Both cases of Theorem 4.3 hold both in our settled model where L, the society's ordering of the candidates, is part of the input and in the model where L is not given and the question is whether there is any societal ordering of the candidates that allows the manipulators to succeed.

We claim that there exists a partition (i.e., a set $A \subseteq I = \{1, 2, \ldots, n\}$ such that $\sum_{i \in A} k_i = \sum_{i \in I - A} k_i$) if and only if the manipulators can cast single-peaked votes (with respect to L) that make p a winner of $(C, S \cup T)$.

From left to right, suppose there is a partition $A \subseteq I = \{1, 2, \ldots, n\}$ with $\sum_{i \in A} k_i = \sum_{i \in I - A} k_i$. Set those voters in T whose weights are k_i, $i \in A$, to $p > a > b$, and set the remaining voters in T (i.e., those with weights k_i, $i \in I - A$) to $p > b > a$. Note that T is single-peaked with respect to L. We have $score_{S \cup T}(p) = 10K + 6K = 16K$ and $score_{S \cup T}(a) = score_{S \cup T}(b) = 15K + K = 16K$, so p is a winner of $(C, S \cup T)$.

From right to left, suppose T is set such that it is single-peaked with respect to L and p is a winner of $(C, S \cup T)$. Without loss of generality, we may assume that p is first in every voter in T. Thus $score_{S \cup T}(p) = 16K$. So, for p to be a winner, we must have $score_T(a) \leq K$ and $score_T(b) \leq K$. But then T corresponds to a partition.

Changing the weights of the voters in S from $5K$ to $5K - 1$ handles the unique-winner case for $\alpha = (3, 1, 0)$.

The scoring protocol $\alpha = (3, 2, 1, 0)$. We again give a reduction from PARTITION to show NP-hardness. We start with the nonunique-winner case.

Given an input to the PARTITION problem, i.e., a set $\{k_1, \ldots, k_n\}$ of n distinct positive integers that sum to $2K$, we form the following instance of Constructive Coalition Weighted Manipulation. The candidate set is $C = \{a, b, c, p\}$ with p being our distinguished candidate, the single-peakedness order L is given by $a \, L \, b \, L \, p \, L \, c$,[20] and the set S of nonmanipulative voters contains two voters, v_1 and v_2. Voter v_1 has weight $11K$ and preference order $c > p > b > a$ and voter v_2 has weight $7K$ and preference order $b > a > p > c$. There are n manipulators in the manipulator set T and the ith manipulator, $1 \leq i \leq n$, has weight k_i. We claim that the manipulators can cast votes that ensure that p is a winner in $(C, S \cup T)$ if and only if $\{k_1, \ldots, k_n\}$ can be partitioned into two subsets that sum to K.

Not counting the manipulators' votes, the candidates have the following scores:

1. $score_S(a) = 14K$,

2. $score_S(b) = 32K$,

3. $score_S(p) = 29K$, and

4. $score_S(c) = 33K$.

Since the total weight of the manipulators is $2K$ and we use scoring vector $(3, 2, 1, 0)$, it is easy to see that irrespective of what votes the manipulators cast, candidate a can never become a winner. Using this fact, it is easy to see that out of the eight preference orders $o_1, o_2, \ldots, o_8$ (see Table 1) that are consistent with L the manipulators can limit themselves to choosing between $o_5 = p > b > a > c$ and $o_7 = p > c > b > a$. (For each of the remaining orders Table 1 explains why either o_5 or o_7 is at least as good an order to cast.)

---

[20]If L is not part of the input, we add two voters to S, namely $a > b > p > c$ and $c > p > b > a$, each of weight 1, who fix the society's order to be our required L (or, equivalently, L's reverse).

	Order	Comments
o_1	$a > b > p > c$	The manipulators are better off casting o_5 instead (p gets more points, a gets fewer points, and the scores of b and c do not change).
o_2	$b > a > p > c$	The manipulators are better off casting o_1 instead (and, in effect, are better off casting o_5) because a cannot become a winner anyway and casting o_1 decreases b's score.
o_3	$b > p > a > c$	The manipulators are better off casting o_5 instead (p gets more points, b gets fewer points, and the scores of a and c do not change).
o_4	$b > p > c > a$	The manipulators are better off casting o_5 instead (p gets more points, b gets fewer points, c gets fewer points, a gets more points but a cannot become a winner anyway).
o_5	$p > b > a > c$	The manipulators might use this order.
o_6	$p > b > c > a$	The manipulators are better off casting o_5 instead (c gets fewer points and a gets more points but a cannot become a winner anyway; the scores of p and b do not change).
o_7	$p > c > b > a$	The manipulators might use this order.
o_8	$c > p > b > a$	The manipulators are better off casting o_7 instead (p gets more points, c gets fewer points, and the scores of a and b do not change).

Table 1: The eight preference orders consistent with $a \, L \, b \, L \, p \, L \, c$ and comments as to which of them the manipulators may choose to use.

Let z_5 be the total weight of the manipulators that vote o_5 and let z_7 be the total weight of the manipulators that vote o_7. Clearly, $z_5 + z_7 = 2K$. Now, including the manipulators, the candidates have the following scores:

1. $score_{S \cup T}(a) = 14K + z_5$,

2. $score_{S \cup T}(b) = 32K + 2z_5 + z_7$,

3. $score_{S \cup T}(p) = 29K + 3z_5 + 3z_7 = 35K$, and

4. $score_{S \cup T}(c) = 33K + 2z_7$.

It is easy to see that p is a winner of election $(C, S \cup T)$ if and only if $32K + 2z_5 + z_7 \leq 35K$ and $33K + 2z_7 \leq 35K$, which is equivalent to

(4.a) $$2z_5 + z_7 \ \leq \ 3K,$$
(4.b) $$2z_7 \ \leq \ 2K.$$

Since $z_5 + z_7 = 2K$, it is easy to see that the above two inequalities hold if and only if $z_5 = z_7 = K$. Thus, the manipulators can ensure that p is a winner if and only if the set $\{k_1, \ldots, k_n\}$ can be partitioned into two subsets that both sum up to K. Since the

reduction is clearly computable in polynomial time, the proof for the nonunique-winner case is complete.

To handle the unique-winner case, we change the weights of the voters in S in the above construction as follows. Voter v_1, whose preference order is $c > p > b > a$, has now weight $11K - 3$, and voter v_2, whose preference order is $b > a > p > c$, has now weight $7K - 2$. Then the candidates have the following scores in (C, S):

1. $score_S(a) = 14K - 4$,

2. $score_S(b) = 32K - 9$,

3. $score_S(p) = 29K - 8$, and

4. $score_S(c) = 33K - 9$.

Regardless of which votes the manipulators cast, a still cannot win, which again leaves $o_5 = p > b > a > c$ (with total weight z_5) and $o_7 = p > c > b > a$ (with total weight z_7) as the only reasonable choices for the voters in T. Now, in election $(C, S \cup T)$, we have the following scores:

1. $score_{S \cup T}(a) = 14K - 4 + z_5$,

2. $score_{S \cup T}(b) = 32K - 9 + 2z_5 + z_7$,

3. $score_{S \cup T}(p) = 29K - 8 + 3z_5 + 3z_7 = 35K - 8$, and

4. $score_{S \cup T}(c) = 33K - 9 + 2z_7$.

Thus, p is a unique winner of election $(C, S \cup T)$ if and only if $32K - 9 + 2z_5 + z_7 < 35K - 8$ and $33K - 9 + 2z_7 < 35K - 8$, which again is equivalent to (4.a) and (4.b). The remainder of the proof is as in the nonunique-winner case. This completes the proof of Theorem 4.3. ❑

In fact, we can extend the idea behind the proofs of part 1 of Theorem 4.1 and part 1 of Theorem 4.3 to obtain the following dichotomy result for 3-candidate scoring protocols in the single-peaked case.

Theorem 4.4 *Consider a 3-candidate scoring protocol, namely, $\alpha = (\alpha_1, \alpha_2, \alpha_3)$, $\alpha_1 \geq \alpha_2 \geq \alpha_3$, $\alpha_1 \in \mathbb{N}$, $\alpha_2 \in \mathbb{N}$, $\alpha_3 \in \mathbb{N}$. For the single-peaked case, the constructive coalition weighted manipulation problem (in both the unique-winner model and the nonunique-winner model) is resistant (indeed, NP-complete) when $\alpha_1 - \alpha_3 > 2(\alpha_2 - \alpha_3) > 0$ and is in P otherwise.*[21]

---

[21] This result holds both in our settled model where L, the society's ordering of the candidates, is part of the input and in the model where L is not given and the question is whether there is any societal ordering of the candidates that allows the manipulators to succeed.

Proof.   Consider a 3-candidate scoring protocol $\alpha = (\alpha_1, \alpha_2, \alpha_3)$, where $\alpha_1 \in \mathbb{N}$, $\alpha_2 \in \mathbb{N}$, $\alpha_3 \in \mathbb{N}$, and $\alpha_1 \geq \alpha_2 \geq \alpha_3$. Without loss of generality, we may assume that $\alpha_3 = 0$. If α_3 were not 0, we could replace our scoring protocol with $(\alpha_1 - \alpha_3, \alpha_2 - \alpha_3, 0)$ (see [HH07]). Suppose that $\alpha_1 - \alpha_3 > 2(\alpha_2 - \alpha_3) > 0$, i.e., $\alpha_1 > 2\alpha_2 > 0$. We show that in this case the constructive coalition weighted manipulation problem is NP-complete. Membership in NP is immediate and to prove NP-hardness we reduce from PARTITION. Given a PARTITION instance, i.e., a set $\{k_1, \ldots, k_n\}$ of n distinct positive integers that sum to $2K$, we construct an instance of the constructive coalition weighted manipulation problem as follows. Let a, b, and p be the candidates in C, where p is the distinguished candidate the manipulators want to make win. There are two nonmanipulators in S, with preferences $a > p > b$ and $b > p > a$ and both having weight $(2\alpha_1 - \alpha_2)K$. There are n manipulators in T, where the ith manipulator has weight $(\alpha_1 - 2\alpha_2)k_i$, $1 \leq i \leq n$. Note that, due to $\alpha_1 > 2\alpha_2 > 0$, each voter has a positive integer weight. We fix the society's order L to be $a \, L \, p \, L \, b$. Note that L and its reverse are the only two orders consistent with the votes in S, so even in the model where L is not part of the input, we will still, in effect, be working with L as the society's order.

We claim that there exists a partition (i.e., a set $A \subseteq I = \{1, 2, \ldots, n\}$ such that $\sum_{i \in A} k_i = \sum_{i \in I - A} k_i$) if and only if the manipulators can cast single-peaked votes (with respect to L) that make p a winner of $(C, S \cup T)$.

Suppose there is a partition, i.e., a set $A \subseteq I = \{1, 2, \ldots, n\}$ such that $\sum_{i \in A} k_i = \sum_{i \in I - A} k_i$. To make p win, set $(\alpha_1 - 2\alpha_2)K$ vote weight of T (say, each of the manipulators corresponding to A) to $p > a > b$ and the remaining manipulators to $p > b > a$. Note that T is single-peaked with respect to L. We have the following scores:

$$
\begin{aligned}
score_{S \cup T}(p) &= 2\alpha_2(2\alpha_1 - \alpha_2)K + 2\alpha_1(\alpha_1 - 2\alpha_2)K = 2\left(\alpha_1^2 - \alpha_2^2\right)K; \\
score_{S \cup T}(a) &= \alpha_1(2\alpha_1 - \alpha_2)K + \alpha_2(\alpha_1 - 2\alpha_2)K = 2\left(\alpha_1^2 - \alpha_2^2\right)K; \\
score_{S \cup T}(b) &= \alpha_1(2\alpha_1 - \alpha_2)K + \alpha_2(\alpha_1 - 2\alpha_2)K = 2\left(\alpha_1^2 - \alpha_2^2\right)K.
\end{aligned}
$$

So p is a winner of election $(C, S \cup T)$.

Conversely, suppose the votes in T can be cast such that p is a winner of $(C, S \cup T)$. Note that

$$
\sum_{c \in C} score_{S \cup T}(c) = 6\left(\alpha_1^2 - \alpha_2^2\right)K.
$$

So, in order for p to be a winner, p's score in $(C, S \cup T)$ needs to be at least $2\left(\alpha_1^2 - \alpha_2^2\right)K$. This can happen only if p is ranked first by every voter in T. And if we set T's votes like that, p's score is exactly $2\left(\alpha_1^2 - \alpha_2^2\right)K$. In order for p to be a winner, a's score and b's score need to be $2\left(\alpha_1^2 - \alpha_2^2\right)K$ each. This means that the vote weight of the voters in T voting $p > a > b$ is equal to the vote weight of the voters in T voting $p > b > a$. Thus there exists a partition.

Since the reduction is polynomial-time computable, the constructive coalition weighted manipulation problem for the scoring protocol $\alpha = (\alpha_1, \alpha_2, 0)$ is NP-complete, provided that $\alpha_1 > 2\alpha_2 > 0$.

The unique-winner case can be handled by multiplying all weights in the above construction by a suitable constant (depending on α_1 and α_2) and then subtracting 1 from the weights in S. The idea is that subtracting 1 from the weights in S ensures that p can become the unique winner if there is a partition, and multiplying the weights, prior to the subtraction, ensures that the subtraction does not have any side effects that would invalidate the (analogue of the) above correctness argument in this modified reduction.

Now suppose $\alpha_1 \geq \alpha_2 \geq 0$ but it is *not* the case that $\alpha_1 > 2\alpha_2 > 0$. We show that in this case the constructive coalition weighted manipulation problem is in P. If $\alpha_2 = 0$ then our scoring protocol is equivalent to plurality for three candidates, and we accept if and only if p is a winner (a unique winner) when all the manipulators rank p first. Let us now focus on the case where $\alpha_2 > 0$, and so, by assumption, $\alpha_1 \leq 2\alpha_2$. Without loss of generality, it is enough to consider the following two orderings L witnessing the society's single-peakedness. (Again, if L is not part of the input, we simply try all societal orderings.) We'll handle the nonunique-winner case and the unique-winner case in parallel.

Case 1: $a\,L\,b\,L\,p$. In this case, if the goal is to make p a winner (a unique winner), a best vote for all manipulators in T is clearly $p > b > a$.

Case 2: $a\,L\,p\,L\,b$. To be consistent with L, the votes allowed in S are $p > a > b$, $p > b > a$, $a > p > b$, and $b > p > a$ (see the proof of part 1 of Theorem 4.1). Due to which votes are allowed in S, we have $2 \cdot score_S(p) \geq score_S(a) + score_S(b)$. It follows that $score_S(p) \geq score_S(a)$ or $score_S(p) \geq score_S(b)$. If we are in the nonunique-winner model, set all voters in T to $p > a > b$ if $score_S(p) \geq score_S(a)$, and to $p > b > a$ if $score_S(p) \geq score_S(b)$. If we are in the unique-winner model and $\alpha_1 > \alpha_2$, set all voters in T to $p > b > a$ if $score_S(a) \geq score_S(b)$ and to $p > a > b$ otherwise. The $\alpha_1 = \alpha_2$ case follows from part 2 of Theorem 4.1.

In both cases, we accept if the manipulation action described above was successful (i.e., made p a winner in the nonunique-winner case (a unique-winner in the unique-winner case) in election $(C, S \cup T)$), and we reject otherwise. ❑

Finally, we mention that despite the NP-hardness results of Theorems 4.3 and 4.4, in all these cases there are polynomial-time manipulation algorithms for the case when the candidate the coalition wants to win is either the top or the bottom candidate in society's input linear order on the candidates.

5   Related Work

The paper that inspired our work is Walsh's "Uncertainty in Preference Elicitation and Aggregation" [Wal07]. Among other things, in that paper he raises the issue of manipulation in single-peaked societies. Our paper follows his model of assuming society's linear ordering of the candidates is given and that manipulative voters must be single-peaked with respect to that ordering. However, our theme and his differ. His manipulation results present cases

where single-peakedness leaves an NP-completeness shield intact. In particular, for both the constructive and the destructive cases, he shows that the coalition weighted manipulation problem for the single transferable vote election rule for three or more candidates remains NP-hard in the single-peaked case. Although our Theorem 4.3 follows this path of seeing where shields remain intact for single-peaked preferences, the central focus of our paper is that single-peaked preferences often remove complexity shields on manipulation and control. Walsh's paper for a different issue—looking at incomplete profiles and asking whether some/all the completions make a candidate a winner—proves both P results and NP-completeness results. We're greatly indebted to his paper for raising and exploring the issue of manipulation for single-peaked electorates.

As mentioned in the main text, Bartholdi and Trick [BT86], Doignon and Falmagne [DF94], and Escoffier, Lang, and Öztürk [ELÖ08] have provided efficient algorithms for testing single-peakedness and producing a valid candidate linear ordering, for the case when votes are linear orders.

Other work is more distant from our work but worth mentioning. Conitzer [Con09] has done an interesting, detailed study showing that (in the model where votes are linear orders) preferences in single-peaked societies can be quickly elicited via comparison queries ("Do you prefer candidate i to candidate j?"). He studies the case when the linear order of society is known and the case when it is not. We mention in passing that we have looked at the issue of preference elicitation in single-peaked societies (where the linear order is given) of approval vectors via approval queries ("Do you approve of candidate i?"). It is immediately obvious that single-peakedness gives no improvement for approval vectors and 1-approval vectors (approval vectors with exactly one 1; this is a vote type and should not be confused with "1-approval" as it would be used in scoring systems, where the actual vote is a linear order). But for j-candidate k-approval vectors (approval vectors with exactly k 1's), $k \geq 1$, it is easy to see that the general-case elicitation query complexity is exactly 0 when $j = k$ and is exactly $j - 1$ when $j > k$, but for the single-peaked case, the elicitation query complexity for ik-candidate k-approval vectors is at most $i - 1 + \lceil \log_2 k \rceil$. That is, we get a savings of a multiplicative factor of about k from single-peakedness.

Single-peaked preferences of course have been studied extensively in political science. We in particular mention that Ballester and Haeringer [BH07] provide a precise mathematical characterization of single-peakedness, that Lepelley [Lep96] shows that single-peakedness removes some negative results about the relationship between scoring protocols and Condorcet-type criteria, and Gailmard, Patty, and Penn [GPP] discuss Arrow's Theorem on single-peaked domains. We refer the interested reader also to the coverage of single-peaked preferences in the excellent textbook by Austen-Smith and Banks [AB04].

6 Conclusions and Future Directions

The central point of this paper is that single-peaked preferences remove many complexity-theoretic shields against control and manipulation. That is, we showed that those shields, already under frequency-of-hardness and approximation attacks from other quarters, for

single-peaked preferences didn't even exist in the first place. It follows that when choosing election systems for electorates one suspects will be single-peaked, one must not rely on results for those systems that were proven in the standard, unrestricted preference model.

This paper's work suggests many directions for future efforts. For single-peaked manipulation of scoring protocols, we gave some NP-complete cases and some P cases. But for manipulation of scoring protocols (in the general model where society is not required to be single-peaked), Hemaspaandra and Hemaspaandra ([HH07], see also [PR07,CSL07]) have provided a dichotomy theorem clearly classifying each case as NP-complete or in P. Can we obtain a dichotomy theorem for manipulation of scoring protocols in single-peaked societies? Theorem 4.4 achieves this for the case of three candidates.

Throughout this paper, single-peaked has meant the unidimensional case. Do the shield removals of this paper hold in, for example, an appropriate two-dimensional (or k-dimensional) analogue? (We mention in passing that every profile of n voters voting by linear orders can be embedded into $\mathbb{R}^n$ in such a way that each voter and candidate is a point in $\mathbb{R}^n$ and each voter prefers c_i to c_j if her Euclidean distance to c_i is less than to c_j.)

In a human society with a large number of voters, even if one issue, e.g., the economy, is almost completely polarizing the society, there are bound to be a few voters whose preferences are shaped by quite different issues, e.g., a given candidate's religion. So it would be natural to ask whether the shield-evaporation results of this paper can be extended even to societies that are "very nearly" single-peaked (see [ELÖ08, Section 6] and [Con09, Section 6] for discussion of nearness to single-peakedness in other contexts).

Finally, we mention that subsequent work of Brandt et al. [BBHH] further explores single-peaked electorates—looking at bribery and at control by partition of voters, and generalizing Theorem 4.4 to each fixed number of candidates.

References

[AB04]   D. Austen-Smith and J. Banks. *Positive Political Theory II: Strategy and Structure.* University of Michigan Press, 2004.

[Bar01]   S. Barberà. An introduction to strategy-proof social choice functions. *Social Choice and Welfare*, 18(4):619–653, 2001.

[BBHH]   F. Brandt, M. Brill, E. Hemaspaandra, and L. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 215–220. AAAI Press. To appear, 2010.

[BEH+] D. Baumeister, G. Erdélyi, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Computational aspects of approval voting. In J. Laslier and R. Sanver, editors, *Handbook of Approval Voting*. Springer. To appear. Available as University of Rochester Computer Science Department Technical Report TR-2009-944, May 2009.

[BFH+08] E. Brelsford, P. Faliszewski, E. Hemaspaandra, H. Schnoor, and I. Schnoor. Approximability of manipulating elections. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, pages 44–49. AAAI Press, July 2008.

[BH07] M. Ballester and G. Haeringer. A characterization of the single-peaked domain, 2007. Manuscript.

[BL76] K. Booth and G. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. *Journal of Computer and System Sciences*, 13(3):335–379, 1976.

[Bla48] D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.

[Bla58] D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.

[BO91] J. Bartholdi, III and J. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

[BT86] J. Bartholdi, III and M. Trick. Stable matching with preferences derived from a psychological model. *Operations Research Letters*, 5(4):165–169, 1986.

[BTT89] J. Bartholdi, III, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[BTT92] J. Bartholdi, III, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modeling*, 16(8/9):27–40, 1992.

[Con09] V. Conitzer. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, 35:161–191, 2009.

[CS06] V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 627–634. AAAI Press, July 2006.

[CSL07] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):1–33, 2007.

[DF94] J. Doignon and J. Falmagne. A polynomial time algorithm for unidimensional unfolding representations. *Journal of Algorithms*, 16(2):218–233, 1994.

[DHO70] O. Davis, M. Hinich, and P. Ordeshook. An expository development of a mathematical model of the electoral process. *American Political Science Review*, 54(2):426–448, 1970.

[ELÖ08] B. Escoffier, J. Lang, and M. Öztürk. Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence*, pages 366–370. IOS Press, July 2008.

[ENR09] G. Erdélyi, M. Nowak, and J. Rothe. Sincere-strategy preference-based approval voting fully resists constructive control and broadly resists destructive control. *Mathematical Logic Quarterly*, 55(4):425–443, 2009.

[ER91] E. Ephrati and J. Rosenschein. The Clarke Tax as a consensus mechanism among automated agents. In *Proceedings of the 9th National Conference on Artificial Intelligence*, pages 173–178. AAAI Press, July 1991.

[ER97] E. Ephrati and J. Rosenschein. A heuristic technique for multi-agent planning. *Annals of Mathematics and Artificial Intelligence*, 20(1–4):13–67, 1997.

[FG65] D. Fulkerson and G. Gross. Incidence matrices and interval graphs. *Pacific Journal of Mathematics*, 15(5):835–855, 1965.

[FHH09] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. How hard is bribery in elections? *Journal of Artificial Intelligence Research*, 35:485–532, 2009.

[FHHR09a] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research*, 35:275–341, 2009.

[FHHR09b] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. A richer understanding of the complexity of election systems. In S. Ravi and S. Shukla, editors, *Fundamental Problems in Computing: Essays in Honor of Professor Daniel J. Rosenkrantz*, pages 375–406. Springer, 2009.

[FHHR09c] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 118–127. ACM Press, July 2009.

[FKN08] E. Friedgut, G. Kalai, and N. Nisan. Elections can be manipulated often. In *Proceedings of the 49th IEEE Symposium on Foundations of Computer Science*, pages 243–249. IEEE Computer Society, October 2008.

[GJ79] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.

[GPP]   S. Gailmard, J. Patty, and E. Penn. Arrow's theorem on single-peaked domains. In E. Aragonés, C. Beviá, H. Llavador, and N. Schofield, editors, *The Political Economy of Democracy*. Fundación BBV. To appear.

[HH07]   E. Hemaspaandra and L. Hemaspaandra. Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1):73–83, 2007.

[HHR07]   E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence*, 171(5–6):255–285, 2007.

[HHR09]   E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Hybrid elections broaden complexity-theoretic resistance to control. *Mathematical Logic Quarterly*, 55(4):397–424, 2009.

[Kre98]   K. Krehbiel. *Pivotal Politics: A Theory of U.S. Lawmaking*. University of Chicago Press, 1998.

[Lep96]   D. Lepelley. Constant scoring rules, Condorcet criteria, and single-peaked preferences. *Economic Theory*, 7(3):491–500, 1996.

[NW87]   R. Niemi and J. Wright. Voting cycles and the structure of individual preferences. *Social Choice and Welfare*, 4(3):173–183, 1987.

[PHG00]   D. Pennock, E. Horvitz, and C. Giles. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proceedings of the 17th National Conference on Artificial Intelligence*, pages 729–734. AAAI Press, July/August 2000.

[PR97]   K. Poole and H. Rosenthal. *Congress: A Political-Economic History of Roll-Call Voting*. Oxford University Press, 1997.

[PR07]   A. Procaccia and J. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 28:157–181, 2007.

[PS99]   B. Peleg and P. Sudhölter. Single-peakedness and coalition-proofness. *Review of Economic Design*, 4(4):381–387, 1999.

[Tri89]   M. Trick. Recognizing single-peaked preferences on a tree. *Mathematical Social Sciences*, 17(3):329–334, 1989.

[Wal07]   T. Walsh. Uncertainty in preference elicitation and aggregation. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 3–8. AAAI Press, July 2007.

[XC08a]   L. Xia and V. Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 109–118. ACM Press, July 2008.

[XC08b]   L. Xia and V. Conitzer. A sufficient condition for voting rules to be frequently manipulable. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 99–108. ACM Press, July 2008.